

## LA-UR-17-27041

Approved for public release; distribution is unlimited.

Title: Understanding Congestion on Omni-Path Fabrics Presentation

Author(s): Gillespie, Lauren Estelle  
Leap, Christopher Aragon  
Cassidy, Daniel Joseph

Intended for: HPC Mini showcase

Issued: 2017-08-07

---

**Disclaimer:**

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

# Understanding Congestion on Omni-Path Fabrics

**Dan Cassidy**  
Indiana University South Bend  
dcassidy@lanl.gov

**Lauren Gillespie**  
Southwestern University  
gillespl@southwestern.edu

**Christopher Leap**  
University of New Mexico  
cleap@unm.edu

LA-UR-17...




- High Performance Clusters → High-Speed Interconnects
- Improved Computational power warrants better networks
- InfiniBand, Gemini, Aries, and recently **OPA**
- However:
  - Large clusters → Heavy network traffic → **Congestion**
- Understand effect of congestion on cluster performance

# Omni-Path Architecture (OPA)

- New high-speed architecture from Intel
- Up to 100 Gb/s transfer rate
- Supports IPoIB, verbs, and psm2 transport protocol
- Performance Counters – for monitoring network performance
  - Transmit Discard Counter
    - Source port down + Malformed Packets + Congestion Discards
    - No differentiation in discard types
    - Included in InfiniBand's counters
  - Congestion Discard Counter
    - Specifically differentiates drops due to congestion from other network errors
    - Provides finer granularity on congestion affecting the network

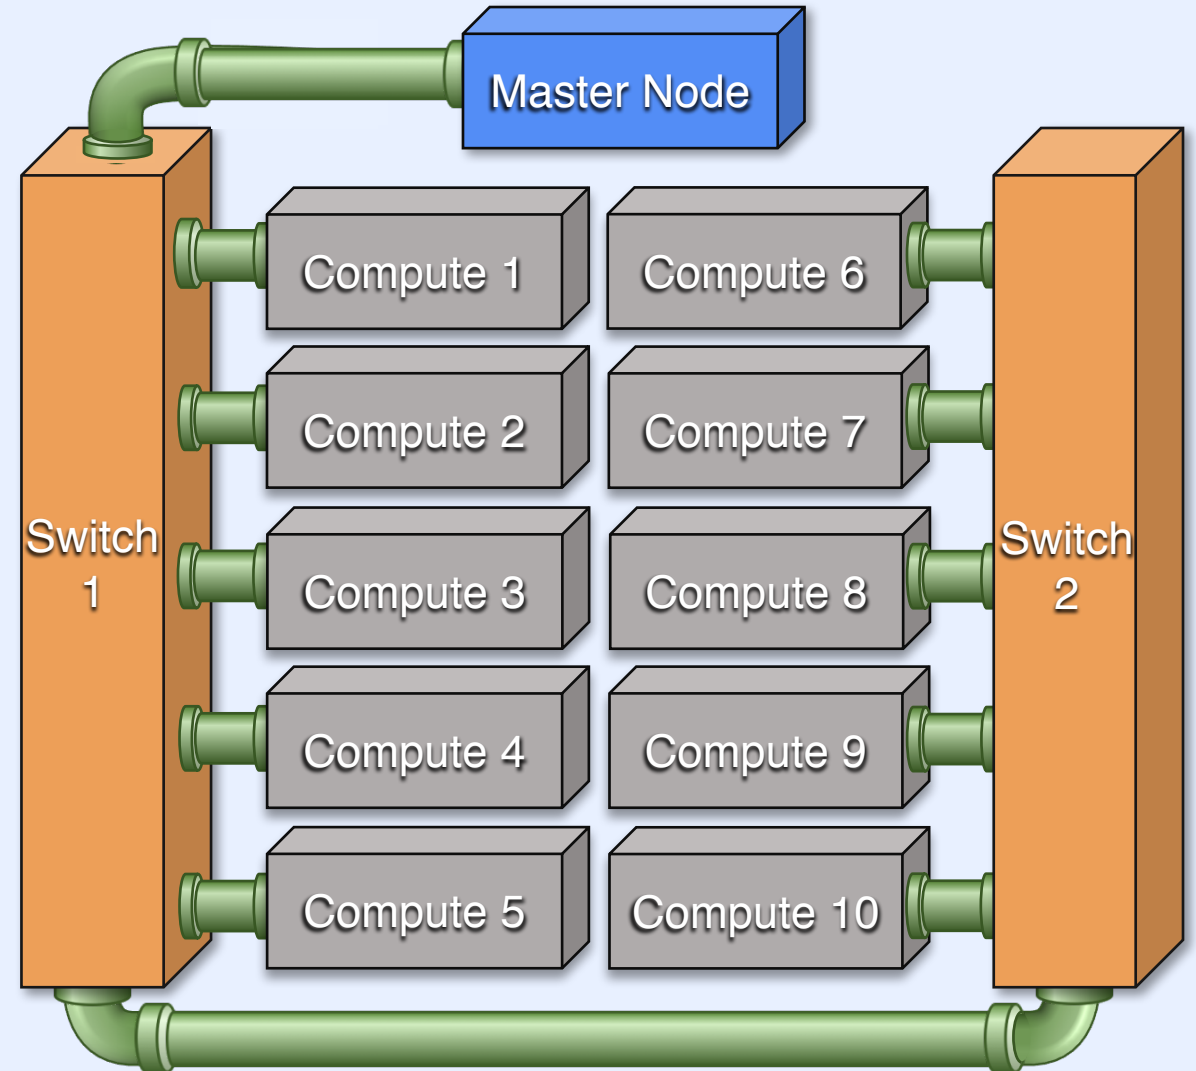
# Congestion Discards

- Counts discards due to three settings:
  - Head-of-queue Lifetime (HoqLife) – time for Hoq packet to enter fabric (in ms)
  - Switch Lifetime (SwitchLife) – time for queued packet to enter fabric (in ms)
  - Stalled Virtual Lanes Count (VLStallCount) – number of consecutive congestion discards allowed before flushing the queue



# System Setup

- Two OPA switches
  - One interlink
- Five compute nodes per switch
- Head node attached to switch 1
- For testing:
  - Each child node pairs to another node on opposing switch
  - $1 \leftrightarrow 6$ ,  $2 \leftrightarrow 7$ , etc.
- Set up designed to force congestion across single interlink




# Experimental Setup

- Compared congestion discards to bandwidth performance
- Bandwidth benchmark: OSU Bi-Directional Bandwidth
- Background noise: creates heavy traffic on network
  - `ib_verbs` test: floods network with queue pairs traffic
  - `dd` 29 GB file: floods network with IP over IB traffic
    - Each node pair transfers 29 GB files across the interlink
- Varied `HopLife` and `SwitchLifetime` separately between tests
- Ran each test for 50 runs




# Why That Didn't Work...

- The Intel P-State driver caused major variance in all results
  - The manually set CPU frequency was ignored by the P-State driver
  - Changed to ACPI driver, fixed OSU baseline variance




Bandwidth baseline with P-State




Bandwidth baseline with ACPI

# Why That Didn't Work Either...

- Furthermore, dd also was causing high data variance
  - Led to consistent discards but inconsistent performance
  - Not an overall uniform data stream, too varied for background noise




Baseline Bandwidth with dd + verbs




Baseline Bandwidth without dd + verbs

# Why it Still Didn't Work...

- Using only verbs for background noise created no discards
  - Led to consistent performance but no discards
  - Psm2 and queue pairs interaction led to bottleneck on NIC



Verbs and OSU bandwidth benchmark




dd and OSU bandwidth benchmark


# >Revised< Experimental Setup

- Compare discards to bandwidth and latency performance
- Varied HoqLife and SwitchLifetime separately between tests
- Also compared IPoIB performance vs. Verbs performance
- Background noise: dd 29 GB file across each node pair
  - dd caused more discards with an acceptable variance
- Bandwidth benchmark: OSU Bi-Directional Bandwidth
- Latency benchmark: OSU Alltoall Latency Test
- Ran each test for 50 runs

# HoqLife Performance Results




dd and OSU latency benchmark with HoqLife




dd and OSU bandwidth benchmark with HoqLife

# SwitchLife Performance Results




dd and OSU latency benchmark with SwitchLifetime




dd and OSU bandwidth benchmark with SwitchLifetime

# IPoIB vs. Verbs Results

- Queue pairs and psm2 leads to much lower performance



Verbs and OSU bandwidth test suite



dd and OSU bandwidth test suite

# Conclusion

- Congestion discards alone are not indicative of network performance degradation
  - However, verbs + psm2 leads to drastically worse performance but no discards
- Data shows weak correlation between congestion discards and performance when HoqLife and SwitchLifetime are enabled
- Performance did not change markedly from baselines
  - Only exception HoqLife and SwitchLifetime 0 (unlimited timer)
- Had higher than preferred variance while using dd



# Future Work

- Larger cluster to experiment with different congestion patterns
  - 11 node cluster unable to force reliable congestion with workload
- Modify VLStallCount parameter, see effects on performance
  - More aggressive congestion parameter, leads to more drops
- Further investigate queue pair interaction with psm2
  - Understand SDMA effects on performance
  - Experienced dip at 16k seen during some tests

# Acknowledgements

**Mentors:** Susan Coulter, Jesse Martinez, and Howard Pritchard

**Program Coordinators:** Amanda Bonnie, Alfred Torrez

**Instructor:** Joan Lucas

**Teaching Assistant:** Hunter Easterday

**Intel Support:** Brian McGinnis

LA-UR-17...

# Questions?

**Dan Cassidy**

Indiana University South Bend  
dcassidy@lanl.gov

**Lauren Gillespie**

Southwestern University  
gillespl@southwestern.edu

**Christopher Leap**


University of New Mexico  
cleap@unm.edu

LA-UR-17...


**Come see our poster from 8 am to 1pm  
Thursday August 10th at the Library!**

# Auxiliary Slides

# High Variance in Results




Dd + OSU bandwidth high variance @ SwitchLife 0




Dd and OSU bandwidth high variance @ HoqLife 0

# Latency + Bandwidth HoqLife Results




dd and OSU latency benchmark with SwitchLifetime




dd and OSU bandwidth benchmark with SwitchLifetime

# Latency + Bandwidth SwitchLife Results



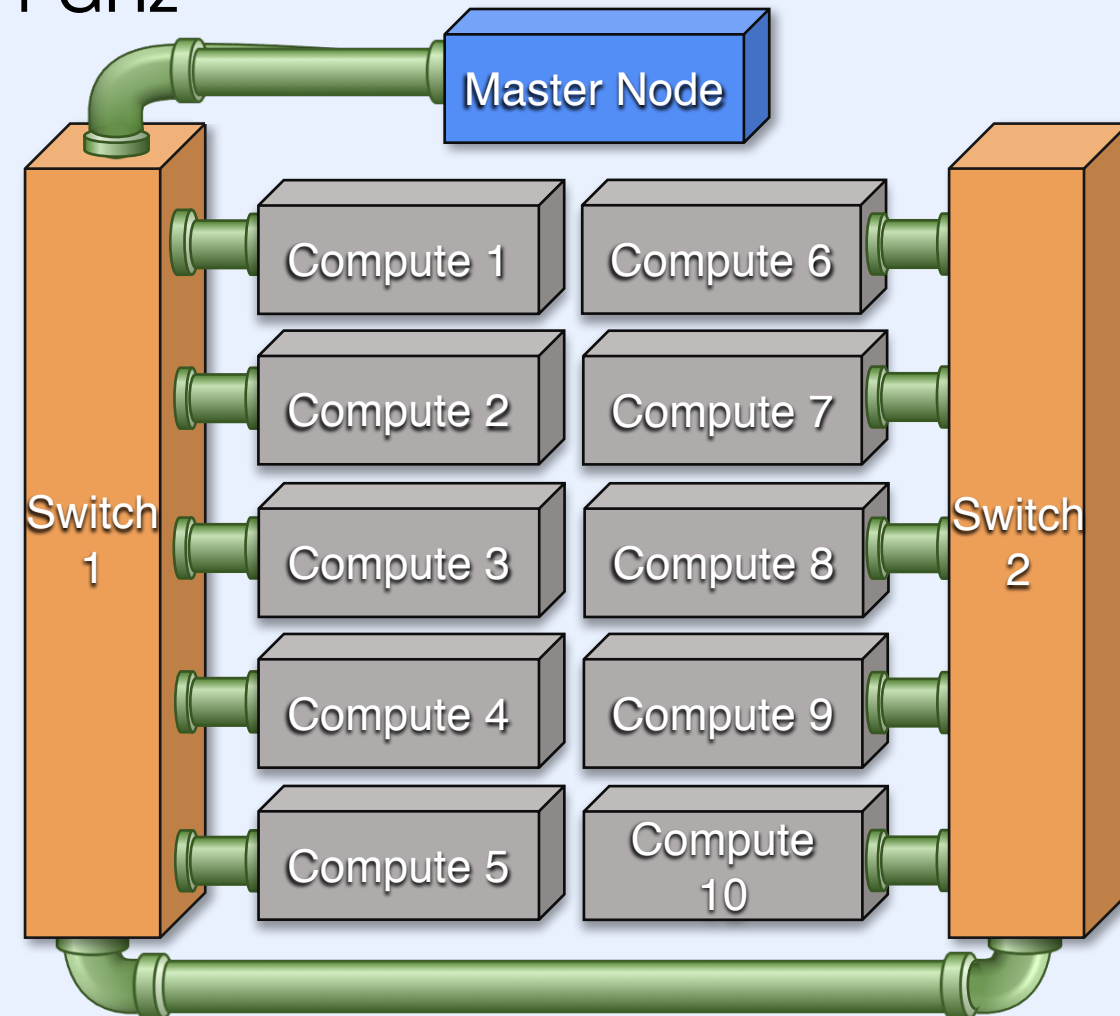
Dd and OSU latency benchmark with HoqLife



dd and OSU bandwidth benchmark with HoqLife

# Node Specifications

- CPU: 2x Intel® Xeon® CPU E5-2620 v4 at 2.1 GHz
- RAM: 64 GB DDR4 at 2133 MHz
- HFI driver: 0.9-294
- OPA firmware: 10.3.0.0.66
- OPA software: 10.4.2.0.7
- OpenMPI: 1.10.4
- OS: CentOS 7.3
- OSU Benchmark: 5.3.2
- PerfTest: 3.0-7
- Verbs Library: Open Fabrics Library





# InfiniBand Vs. Omni-Path

## InfiniBand

- One Counter:
  - Transmit Discards
    - Source port down
      - Output port not active
    - Malformed Packets
      - Packet length exceeded
    - Congestion
      - Head-of-queue Life (HoqLife)
      - Switch Lifetime (SwitchLife)
      - Stalled Count (VLStallCount)


No differentiation in discard types

## Omni-Path


- Two different counters:
  - Transmit Discards
    - Source port down
    - Malformed Packets
    - Congestion
  - Congestion Discards
    - Head-of-queue Life (HoqLife)
    - Switch Lifetime (SwitchLife)
    - Stalled Count (VLStallCount)

Specifically calls out discards due to congestion

# Performance Dip at 16k



OSU Bandwidth + Verbs default settings



OSU Bandwidth + Verbs SwitchLife 60